UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/665,808	09/19/2003	Hao Xu	A033	6637

36378 7590 06/10/2010

VMWARE, INC.
DARRYL SMITH
3401 Hillview Ave.
PALO ALTO, CA 94304

EXAMINER	
RECEK, JASON D	

ART UNIT	PAPER NUMBER
2442	

NOTIFICATION DATE	DELIVERY MODE
06/10/2010	ELECTRONIC

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Notice of the Office communication was sent electronically on above-indicated "Notification Date" to the following e-mail address(es):

ipteam@vmware.com
ipadmin@vmware.com

Office Action Summary	Application No.	Applicant(s)	
	10/665,808	XU ET AL.	
	Examiner	Art Unit	
	JASON RECEK	2442	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE *3* MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1)☒ Responsive to communication(s) filed on *03 February 2010*.

2a)☒ This action is **FINAL**. 2b)☐ This action is non-final.

3)☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

4)☒ Claim(s) *1-11,13,14,30 and 49-85* is/are pending in the application.

4a) Of the above claim(s) _____ is/are withdrawn from consideration.

5)☐ Claim(s) _____ is/are allowed.

6)☒ Claim(s) *1-11, 13-14, 30 and 49-85* is/are rejected.

7)☐ Claim(s) _____ is/are objected to.

8)☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9)☐ The specification is objected to by the Examiner.

10)☐ The drawing(s) filed on _____ is/are: a)☐ accepted or b)☐ objected to by the Examiner.

Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

11)☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

12)☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

a)☐ All b)☐ Some * c)☐ None of:

1.☐ Certified copies of the priority documents have been received.

2.☐ Certified copies of the priority documents have been received in Application No. _____.

3.☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) ☒ Notice of References Cited (PTO-892)
2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
3) ☐ Information Disclosure Statement(s) (PTO/SB/08) Paper No(s)/Mail Date _____.
4) ☐ Interview Summary (PTO-413) Paper No(s)/Mail Date. _____ .
5) ☐ Notice of Informal Patent Application
6) ☐ Other: _____.

DETAILED ACTION

This is in response to the amendment filed on February 3rd 2010.

Status of Claims

Claims 1-11, 13-14, 30 and 49-85 are pending. Claims 71-85 are newly added.

Response to Arguments

1. Applicant's arguments filed 2/3/10 have been fully considered but they are not persuasive. Applicant asserts:

a. None of the references teach of suggest "based on VM-specific information, selecting a NIC from the plurality of NICs" (pg. 15). In support of this argument applicant states that the mechanism of Vega for allocating processor time could not be used to allocate a NIC resource (pg. 16). This argument is not persuasive. Applicant acknowledges a NIC is a computer resources (pg. 16). Vega clearly teaches allocating "resources" of a computer system (col. 3 ln. 35-36). Vega explicitly discloses using VM-specific information when allocating "the allocation of resources may be accomplished according to a proportional weight assigned to each virtual machine" (col. 3 ln. 39-41). In further support of this argument applicant states Vepa does not mention virtual machines (pg. 17). Although factual, this is not persuasive with respect to the argument. Vega was relied upon for teaching VM specific information. Vepa was

merely cited for its disclosure of using source information to select a NIC (i.e. load balancing). Applied to the system of Vega (i.e. when the source is a VM) such a combination would use VM specific information (as disclosed by Vega) for the purpose of selecting a NIC (Vepa). Thus the combination of Mahalingam, Vega and Vepa at least suggests "based on the NIC management information and the VM-specific information, selecting a NIC from the plurality of NICs and transferring the outgoing data from" as recited by the claims.

b. There is no motivation to combine (pg. 18-21). In support of this argument applicant states there is no specific motivation in the prior art (pg. 18). This argument is not persuasive. Although applicant references *KSR*, it seems applicant is saying since there is no teaching, suggestion or motivation, the references cannot be combined. This is not the standard under *KSR*. Applicant further states the references cannot be combined with no change in their respective functions or would require significant modification (pg. 18-19). This argument is also not persuasive. The base reference Mahalingam as modified by the proposed combination of references is still performing load balancing (i.e. NIC selection). Therefore the function has not changed as suggested by applicant. As shown in Fig. 4 of Mahalingam, the drivers have access to the packets. Vepa teaches the source information is located within the packets. Since Mahalingam already has access to this information it would not require significant modification as suggested by applicant. Applicant again asserts the allocation technique of Vega cannot be applied to NICs (pg. 20). This argument is not persuasive for the same reason, Vega teaches allocating resources, applicant acknowledged a NIC

is a resource (see a, above). Applicant asserts (pg. 21) that the only motivation to use VM-specific information to select NICs would be for QoS and this could have only come from their own disclosure (i.e. hindsight). This is not persuasive. Vega explicitly teaches QoS as allocating resources to ensure a minimum level (col. 4 ln. 9-11). Thus applicant's suggestion that this motivation could have only come from improper hindsight is not persuasive. Furthermore, the rejection provided motivation to combine, specifically it was presented that one of ordinary skill in the art would understand a NIC to be a “resource” and therefore applying the allocation scheme which uses VM specific information to a NIC is merely the combination of known elements. Vepa provides explicit suggestion because it teaches selecting a NIC using source data. VMs were known in the art at the time of the invention (as evidenced by Vega). Therefore this is also the combination of known elements.

c. Macchiano does not teach the limitations of claims 30 and 64 because it teaches virtual NICs and not physical NICs (pg. 22). This is not persuasive. The last office action clearly indicated that Macchiano did not disclose physical NICs, Mahalingam was relied upon for teaching "physical NICs". Applicant also states no disclosure was found in Macchiano concerning the "determining ..." limitation (pg. 22). This portion of the rejection has been expanded upon below to help applicant understand how Macchiano teaches this step. Applicant further states that Macchiano does not disclose “transferring the outgoing data frame … or discarding …”. This is not persuasive. The office action did not rely upon Macchiano for teaching these features. Agreeing with the examiner will not over come a rejection. Finally, applicant’s

conclusory statement (pg. 22) that “none of the remaining cited references teach this feature” is a merely allegation of patentability and is not persuasive. Applicant's arguments fail to comply with 37 CFR 1.111(b) because they amount to a general allegation that the claims define a patentable invention without specifically pointing out how the language of the claims patentably distinguishes them from the references.

d. Dependent claims are allowable for similar reasons (pg. 23-24). This is not persuasive because the arguments with respect to the independent claims were not persuasive. Applicant also discusses the Ishizaki reference but it is not entirely clear what applicant’s argument is (pg. 23). It seems applicant is suggesting that Ishizaki does not teach discarding outgoing data if a decision is made not to transfer as recited by claims 67 and 69. Applicant states that transferred does not mean outgoing and Ishizaki only discards incoming packets (pg. 23). This argument is not persuasive. Ishizaki is not only concerned with incoming packets as suggested by applicant, it clearly teaches the purpose of the invention is to “transfer” packets (col. 2 ln. 30-35). Thus when Ishizaki makes a decision not to transfer a packet and discards the packet (Fig. 13) it discards an outgoing data frame “if a decision is made not to transfer the outgoing data frame” as recited by the claims. It seems applicant may be narrowly interpreting the limitation "outgoing data frame". Applicant is reminded that claims are read with their broadest reasonable interpretation in light of the specification but limitations from the specification are not read into the claims.

e. New claims 71-85 are allowable for similar reasons as claim 1 (pg. 24). As discussed above claim 1 is not allowable. Therefore this argument is not persuasive.

Claim Rejections - 35 USC § 103

2. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

> (a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negatived by the manner in which the invention was made.

3. Claims 1-10, 51-60, 68, 70-79, 82 and 85 are rejected under 35 U.S.C. 103(a) as being unpatentable over Mahalingam et al. US Pat. No. 6,208,616 B1 in view of Vega US Pat. No. 7,136,800 B1 and Vepa et al. US 6,567,377 B1.

Regarding claim 1, Mahalingam discloses “a method for responding to a request to transfer an outgoing data frame” as a system for re-routing network packets in a computer (col. 3 ln. 49-50), "the outgoing data frame comprising at least data to be transmitted and at least one of a layer 2 and layer 3 destination address” a MAC packet (layer 2) with address (Fig. 8), “virtualization software comprising one or more layers of software" virtual adapter (col. 10 ln. 54-58, "a plurality of physical network interface cards (NICs)" a computer system with a first NIC and a second NIC (col. 2 ln. 45-49).

Mahalingam also discloses:

"obtaining access by a NIC manager to the outgoing data frame" manages data frames for transfer (Fig. 9-10), "the NIC manager being a component of the virtualization software" virtual adapter (Fig. 4) is software, "receiving, in the NIC manager, NIC management information" (col. 3 ln. 3-10), "management information related to one or more of the plurality of NICs" detecting failures on a NIC, such error information is management information that relates to one or more of the NICs (col. 3 ln. 59-66), and "based on the NIC management information [...] selecting a NIC from the plurality of NICs and transferring the outgoing data frame to the computer network over the selected NIC" as a system that can perform load sharing of packets across a plurality of NICs by using NIC loads or error information as a factor and switching between NICs if one fails, thereby transferring over the selected NIC (see abstract, col. 15 ln. 25-35, and Fig. 10).

Mahalingam does not explicitly teach "virtual computer system comprising one or more virtual machines (VMs)", "the outgoing data frame being provided by one of the VMs" or receiving and using VM-specific information in the decision making process. However Vega teaches making a decision "based on [...] the VM-specific information" by allocating resources among multiple virtual machines running on a physical computer. Vega explicitly teaches using VM-specific information (col. 3 ln. 65—col. 4 ln. 7) to manage the host computer's resources. Although Vega is directed to allocating processor time, one skilled in the art understands that a NIC is a simply a computer resource and that similar allocation methods can be used. It would have been obvious to one of ordinary skill in the art at the time of the invention to modify Mahalingam with

the VM resource allocation technique taught by Vega for the purpose resource allocation. Vega suggests that by using VM specific data requirements may be met (col. 3 ln. 64-65). Vega also discloses one or more virtual machines (Fig. 2) and "outgoing data frame being provided by one of the VMs" as the VMs pass data to the hardware, thus they provide the data frame (Fig. 2).

Motivation is also provided by Vepa which teaches a load balancing scheme for selecting physical NICs that takes into consideration the source of the data (col. 3 ln. 40 – col. 4 ln. 17), applied to this case the source represents the virtual machine. Given these teachings, one of ordinary skill in the art at the time of the invention would understand that it may be advantageous to use NIC and sender information (VM specific information) to select a physical network interface for the purpose of load balancing. This combination is merely the combination of known elements (VM specific information – Vega, NIC load balancing – Mahalingam) according to their established functions in order to yield a predictable result.

Regarding claim 2, the additional limitation, "in which the VM-specific information indicates an amount of network bandwidth that is allocated to a VM that requested the data transfer" is suggested by Vega as apportioning a percentage of resources to each virtual machine (see paragraph 9, lines 17-18), or in the alternative an absolute capacity (see paragraph 11, lines 1-2). The motivation to combine the two references was set out in the rejection of claim 1.

Regarding claim 3 the additional limitation, "decision is made not to transfer the data because transferring the data would cause the VM's allocation of network bandwidth to be exceeded" is suggested by Vega. Vega teaches that if a VM were to exceed its allocation of resources, the operation would not be allowed (see paragraph 17, lines 11-15). Motivation to combine is the same rationale as used in claim 1 rejection.

Regarding claim 4 the limitation, "in which the VM-specific information indicates the priority of the VM that requested the data transfer relative to the priorities of other virtual machines" is taught by Vega (see paragraph 11, lines 12-15) where priorities are assigned to VMs for the purpose of resource allocation. The motivation to combine these references follows the same rationale as used in claim 1 rejection.

Regarding claim 5 the limitation, "in which the NIC management information indicates which one or more of the plurality of NICs is available for the transfer of data" is disclosed by Mahalingam. The system in Mahalingam controls which NIC to use, to do this it is inherent that a list of available NICs is kept (see Fig. 2, steps 52-66). Motivation to combine is the same as that used in the claim 1 rejection.

Regarding claim 6, Mahalingam discloses the additional limitation "in which the NIC management information further indicates a pending data transfer load for each of the available NICs" as a system that chooses which NICs to use based on an algorithm

that includes load information (see column 15, lines 30-35). Motivation to combine is the same as that used in the claim 1 rejection.

Regarding claim 7, Mahalingam discloses "in which a load distribution function, based on the NIC management information [...] is used in selecting a NIC over which to transfer the data" as a system that chooses a NIC based on an algorithm that will choose a NIC that is less loaded than another NIC (see column 15, lines 30-35). The motivation to combine Mahalingam and Vega is the same as stated in the claim 1 rejection.

Regarding claim 8, Mahalingam discloses "the ... data transfer requests are routed over the second NIC if the first NIC is not available" as a system having a primary and secondary NIC where traffic is directed to the primary NIC until it fails and thereafter traffic is directed to the remaining NIC (col. 5 ln. 40-52), and "the ... data transfer requests are routed over the first NIC if the second NIC is not available" as analyzing all NICs for failure and routing data accordingly (col. 5 ln. 58-62, Fig. 2).

Mahalingam does not explicitly disclose "a first VM's data transfer" however as discussed in claim 1, Vega teaches allocating resources among virtual machines (col. 3 ln. 35-40).

Although Mahalingam and Vega do not explicitly disclose, "in which a first VM's data transfer requests are substantially always routed over a first NIC as long as the first NIC is available, and a second VM's data transfer requests are substantially always

routed over a second NIC as long as the second NIC is available" this concept is well known in the art (see response to arguments) and yields predictable results. It would have been obvious to one of ordinary skill in the art at the time of the invention to associate a virtual machine with a NIC and use that NIC to transfer data to and from that virtual machine.

Regarding claim 9, Mahalingam and Vega do not disclose "in which the first VM's data transfer requests are distinguished from the second VM's data transfer requests by reference to a source physical address contained in a header of each data transfer request", however Vepa teaches this as a system that balances loads over multiple network interfaces (abstract). Vepa describes a system where an application can send and receive data specific to it (Fig. 3, col. 3), thus the data transfer request inherently must include an address to receive a response (source port, col. 3 ln. 60-65), this address allows one to distinguish between different data originators (VMs). It would have been obvious to one of ordinary skill in the art at the time the invention was made to combine Vepa because a return address to distinguish the sender is a necessity for any type of two-way communication.

Regarding claim 10, Mahalingam discloses "in which the management information indicates whether a failover is occurring on one of the NICs" as a system that detects NIC failures and determines which NIC to use based on this information

(see column 4, lines 30-40). The motivation to combine Mahalingam and Vega is stated above.

Regarding claims 51-60, Applicant states these are susbstantively comparable to the original claims, as amended. Therefore, they are rejected for the same reasons.

Specifically, claims 51-60 correspond to claims 1-10. Therefore, they are rejected for similar reasons.

Regarding claim 68, Mahalingam discloses a “NIC manager” that has access to data besides the VM data as managing a NIC that is connected to a protocol stack, thus all data will be sent through it (col. 2 ln. ln. 45-52).

Regarding claim 70 it corresponds to claim 68 and thus is rejected for similar reasons.

Regarding claim 71, it substantially corresponds to claim 1. The corresponding parts are rejected for similar reasons. Mahalingam does not explicitly disclose "VM-specific information being at least one of an identity of the one VM ... a priority ... or an amount of bandwidth" however this is taught by Vega as a proportional weight (col. 3 ln. 40). The motivation to combine is the same as that given above.

Claims 72-79 correspond to claims 2 and 4-10 respectively, therefore they are rejected for similar reasons.

Claim 82 corresponds to claim 3 and thus is rejected for similar reasons.

Claim 85 corresponds to claim 68 and thus is rejected for similar reasons.

4. Claims 67, 69 and 81 are rejected under 35 U.S.C. 103(a) as being unpatentable over Mahalingam, Vega and Vepa as applied to claims 1 and 51 above, and further in view of Ishizaki et al. US 6,810,421 B1.

Regarding claim 67, Mahalingam, Vega and Vepa do not explicitly disclose "discarding the outgoing data frame" however this is taught by Ishizaki as discarding an outgoing packet (col. 8 ln. 42-50). It would have been obvious to one of ordinary skill in the art at the time of the invention to modify Mahalingam with the discard feature taught by Ishizaki. Ishizaki suggests that discarding packets can reduce administration overhead (col. 2 ln. 5-63).

Regarding claim 69, it corresponds to claim 67 and thus is rejected for similar reasons.

Claim 81 corresponds to claim 67 and thus is rejected for similar reasons.

5. Claims 11, 14, 61, 80 and 84 are rejected under 35 U.S.C. 103(a) as being unpatentable over Mahalingam in view of Vega and Vepa and further in view of Rietschote et al. U.S. Pat. No. 7,203,944 B1.

Regarding claim 11, Mahalingam and Vega do not disclose, "in which the VM that has requested the data transfer is temporarily suspended if a failover is occurring on one of the NICs". However Rietschote does teach suspending a virtual machine to balance load (see col. 1 lines 8-10 and col. 7 lines 4-5). The motivation to combine is load balancing which is apparent from the title of the invention. Thus it would have been obvious to one of ordinary skill in the art at the time the invention was made to suspend the VM when a NIC was failing, this is simply a way of load balancing.

Regarding claim 14, the combination of Mahalingam, Vega and Vepa does not explicitly disclose, "a further decision is made whether to migrate the VM that requested the data transfer to another computer system" however this is taught by Rietschote as a system for performing load balancing by migrating VMs from one computer system to another (see paragraph 21). The motivation for combining Rietschote is similar to the motivation set out in the claim 11 rejection.

Claim 61 corresponds to claim 11 and thus is rejected for similar reasons.

Regarding claim 80, it corresponds to claim 11 and thus is rejected for similar reasons.

Regarding claim 84, it corresponds to claim 14 and thus is rejected for similar reasons.

6. Claims 13, 62-63 and 83 are rejected under 35 U.S.C. 103(a) as being unpatentable over the combination of Mahalingam, Vega, Vepa and Ishizaki as applied to claim 67 above, and further in view of Rietschote et al. US 7,203,944 B1.

Regarding claim 13, the combination of Mahalingam, Vega, Vepa and Ishizaki does not explicitly disclose "if a decision is made not to transfer the data, a further decision is made whether to suspend the VM that requested the data transfer" however this is taught by Rietschote as a system that suspends VMs to balance the load. In the present invention it is assumed that when a decision is made not to transfer this is because the VM is exceeding its share of resources, thus an act of load balancing needs to occur. Rietschote teaches suspending the VM as a way of load balancing (see paragraph 24). The motivation to combine the load balancing features of Rietschote was presented in the rejection of claim 11 above.

Regarding claims 62-63, Applicant states these are susbstantively comparable to the original claims, as amended. Therefore, they are rejected for the same reasons.

Specifically, claims 62-63 correspond to claims 13-14.

Regarding claim 83, it corresponds to claim 13, thus it is rejected for similar reasons.

7. Claims 30 and 64 are rejected under 35 U.S.C. 103(a) as being unpatentable over Macchiano U.S. Pat. 7,111,303 B2 in view of Vega U.S. 7,136,800 B1 and in further view of Mahalingam et al. U.S. Pat. 6,208,616 B1.

Regarding claim 30, Macchiano discloses, "a method for responding to requests to transfer data from a virtual computer system and to a physical computer network" as a way for users on a virtual machine to communicate using Internet Protocol (see column 3, lines 50-52). Macchiano further discloses, "the virtual computer system comprising a first VM and a second VM" as a virtual machine operating system having a first and second user portion (see column 3, lines 53-54, Fig. 1 components 12, 14; col. 4 lines 50-54). Macchiano also discloses, "the virtual computer system also comprising a first … network interface card (NIC) and a second … NIC for connecting to the computer network" as describing each user portion having a virtual NIC, thus the system comprises a first NIC and a second NIC (see col. 3 lines 56-58 and Fig. 1 comp. 42, 44; col. 5 lines 4-6).

Macchiano further discloses, "for each data transfer request: determining which VM within the virtual computer system is involved in the requested data transfer; and if

the first VM is involved in the requested data transfer, transferring the data over the first NIC" first VM application passes data to first NIC this is performed via a base portion maintaining a table of IP addresses by which the device driver addresses its respective NIC (the table of addresses allows for determining which VM is involved in transfer), and where the IP datagram from the first user portion is passed to the first NIC (see column 3, lines 60-66), determination functionality (col. 4 ln. 1-5).

Macchiano does not explicitly disclose "determining that the first VM has a higher priority than the second VM" however this is taught by Vega as assigning a proportional share (priority) to a virtual machine (col. 3 ln. 39-41). It would have been obvious to one of ordinary skill in the art at the time of the invention to modify Macchiano with the priority factor taught by Vega for the purpose of allocating resources. Using priority to give preferential treatment to an application or process is well known in the art and yields predictable results (as evidenced by Vega).

The combination of Macchiano and Vega does not explicitly disclose "physical NICs" or "determining that the second NIC is not available for transferring data" or "in response to determining that the second NIC is not available, discarding the data" however this is taught by Mahalingam et al. as multiple physical NICs (Fig. 1), detecting failure of a NIC (col. 5 ln. 44-48, Fig. 2) and discarding data related to a secondary NIC (col. 15 ln. 12-18). More specifically, Mahalingam teaches discarding packets received from an adapter that has failed (i.e. not available), or in other words, it teaches only passing packets from NICs that are "in use" (col. 11 ln. 15-31). It would have been obvious to one of ordinary skill in the art at the time of the invention to modify

Macchiano with the error checking taught by Mahalingam for the purpose of transferring data. Error checking is well known in the art and yields predictable results (as evidenced by Mahalingam).

Claim 64 is a medium claim that corresponds to the method of claim 30, therefore it is rejected for similar reasons.

8. Claims 49-50 are rejected under 35 U.S.C. 103(a) as being unpatentable over the combination of Macchiano, Vega and Mahalingam as applied to claim 30 above, and further in view of Rietschote.

Regarding claims 49-50, Applicant states these are susbstantively comparable to the original claims, as amended. Therefore, they are rejected for the same reasons.

Specifically, claims 49-50 correspond to claims 13-14.

9. Claims 65-66 are rejected under 35 U.S.C. 103(a) as being unpatentable over the combination of Macchiano, Vega and Mahalingam as applied to claim 64 above, and further in view of Rietschote et al. U.S. Pat. No. 7,203,944 B1.

Claims 65-66 correspond to claims 13-14, therefore they are rejected for similar reasons.

Conclusion

10. The prior art made of record and not relied upon is considered pertinent to applicant's disclosure.

Kitai et al. US 5,948,069 discloses selecting a communication path for quality of service (abstract).

Donovan et al. US 2004/0221285 A1 discloses virtual machine management including allocation of resources (abstract, paragraph 20).

Oyamada et al. US 6,802,062 B1 discloses migrating virtual machines (abstract).

Flynn, JR US 6,453,392 B1 discloses sharing devices between virtual machines (abstract).

Nelson et al. US 7,424,710 B1 discloses a VM using multiple NICs (abstract, Fig. 3).

11. **THIS ACTION IS MADE FINAL.** Applicant is reminded of the extension of time policy as set forth in 37 CFR 1.136(a).

A shortened statutory period for reply to this final action is set to expire THREE MONTHS from the mailing date of this action. In the event a first reply is filed within TWO MONTHS of the mailing date of this final action and the advisory action is not mailed until after the end of the THREE-MONTH shortened statutory period, then the shortened statutory period will expire on the date the advisory action is mailed, and any extension fee pursuant to 37 CFR 1.136(a) will be calculated from the mailing date of

the advisory action. In no event, however, will the statutory period for reply expire later than SIX MONTHS from the mailing date of this final action.

Any inquiry concerning this communication or earlier communications from the examiner should be directed to JASON RECEK whose telephone number is (571)270-1975. The examiner can normally be reached on Mon - Fri 9:00am-5:30pm.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Philip Lee can be reached on (571) 272-3967. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see http://pair-direct.uspto.gov. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

/Jason Recek/
Examiner, Art Unit 2442
(571) 270-1975

/Philip C Lee/
Acting Supervisory Patent Examiner, Art Unit 2442